

Final Report for AOARD Grant FA2386-10-1-4164 **"Integrating Logical and non-Logical Reasoning"**

**Date: 11/30/2011**

**Name of Principal Investigator(s):**

Maurice Pagnucco  
David Rajaratnam  
Claude Sammut  
Michael Thielscher

- e-mail address: {morri | daver | claudie | mit}@cse.unsw.edu.au
- Institution: The University of New South Wales
- Department: School of Computer Science and Engineering
- Mailing Address:  
School of Computer Science and Engineering  
The University of New South Wales  
Sydney, NSW, 2052  
AUSTRALIA
- Phone: +61-2-9385 5518
- Fax: + 61-2-9385 4071

Period of Performance: 09/13/2010 – 11/13/2011

**Abstract:** Short summary of most important research results: **Why** (the work was done), **what** (was accomplished), and **so what** (how did this push scientific frontiers or advance the field). This will be used for archival purposes and literature searches.

This project sets out to explore the interaction between logical and non-logical reasoning systems. It is primarily concerned with forging links between two relatively distinct sub-fields within artificial intelligence and to provide both theoretical and practical benefits.

Why: The project's fundamental goal was to explore methods for practical reasoning in the context of robotics. Robotics has seen quite significant advances in the recent past. With increases in miniaturization, computational power, sensor technology and the advent of robotic middleware (e.g., the Robot Operating System (ROS) used in this project and increasingly adopted by the major robotics research groups worldwide) we now have quite sophisticated robotic systems. However, there is a gap in higher-level reasoning and coordination of lower level behaviours. This project bridges that gap by integrating higher level logical reasoning techniques with lower level non-logical techniques.

What: The project has made advances in dealing with erroneous information, approximate reasoning, topological reasoning and languages for practical reasoning systems.

So what: These developments are currently being implemented on a domestic robot (a Segway RMP) operating in a home environment. The project therefore demonstrates practical reasoning in complex environments.

**Introduction:** Include a summary of specific aims of the research and describe the importance of this research and what is the ultimate goal.

Advances in robotics have led to sophisticated sensors and actuators as well as software environments that provide for quite advanced capabilities in particular with regard to mapping, localization, navigation, vision and manipulation. However the high-level coordination of these skills or behaviours has not yet reached the same level of sophistication. The ultimate goal of this project is to investigate practical reasoning by integrating non-logical reasoning methods that have been successfully used in low-level behaviours (e.g., for SLAM, navigation, object recognition, etc.) and logical methods for

Report Documentation Page			Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.				
1. REPORT DATE <b>23 FEB 2012</b>		2. REPORT TYPE <b>Final</b>		3. DATES COVERED <b>16-09-2010 to 15-09-2011</b>
4. TITLE AND SUBTITLE <b>Integrating Logical and non-Logical Reasoning</b>			5a. CONTRACT NUMBER <b>FA23861014164</b>	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) <b>Maurice Pagnucco</b>			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>University of New South Wales,UNSW,Sydney NSW 2052,Australia,AU,2052</b>			8. PERFORMING ORGANIZATION REPORT NUMBER <b>N/A</b>	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) <b>AOARD, UNIT 45002, APO, AP, 96338-5002</b>			10. SPONSOR/MONITOR'S ACRONYM(S) <b>AOARD</b>	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) <b>AOARD-104164</b>	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT <b>This project explores the interaction between logical and non-logical reasoning systems. It is primarily concerned with forging links between two relatively distinct sub-fields within artificial intelligence and to provide both theoretical and practical benefits. The project's fundamental goal was to explore methods for practical reasoning in the context of robotics. The project made advances in dealing with erroneous information, approximate reasoning, topological reasoning and languages for practical reasoning systems. These developments are currently being implemented on a domestic robot (a Segway RMP) operating in a home environment. The project therefore demonstrates practical reasoning in complex environments.</b>				
15. SUBJECT TERMS <b>Artificial Intelligence, Computer Science</b>				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES <b>10</b>
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>		

high-level coordination of these behaviours. To this end we have developed theoretical solutions for dealing with erroneous information and for approximate reasoning. We have also developed implementations to facilitate topological reasoning and two languages for high level control – one based on a variant of the cognitive robotic programming language GOLOG developed by one of the investigators on this project and another, ALProlog, originally developed for general-game playing by another of the investigators on this project. This has led to the implementation of a ROS framework for controlling a robot in a home environment which provides the experimental setting for the project's developments.

**Experiment:** Description of the experiment(s) performed and the facilities used to perform the work.

A ROS based framework has been developed to provide a platform for experimenting with the integration of logical reasoning and robot control. The platform allows a logical reasoner to send instructions to, and receive feedback from, a ROS controlled robot. Currently, the framework provides a plug-in for the cognitive robot programming language GOLOG (Levesque *et al.* 1997). We use a GOLOG variant developed by one of the investigators on this project (Levesque and Pagnucco 2000) however this can be easily extended to other reasoners and, in particular, we are currently extending the framework to work with the ALPprolog system (Drescher and Thielscher 2011) developed by one of the investigators on this project.


As the framework is based on ROS, we are able to experiment with robot control both in simulation and using real robots in the laboratory. Simple human operator interaction can be achieved using a basic text-based interface while richer more complex interactions can be provided using the multi-modal MICA/FrameScript system (Kadous and Sammut 2004) developed by one of the investigators on this project. FrameScript allows for the development of sophisticated grammars for conversational agents. While we have not yet incorporated a text-to-speech engine for this project, we have done so in the past using the commercial Dragon engine and plan to do this at a later stage.

A basic set of interactions have been implemented in GOLOG to highlight the types of high-level reasoning tasks that would be expected of a domestic robot. For example, the robot can be instructed to go to a room in the house (see Figure 1) as well as to deliver objects to individuals. When the robot is told to "deliver the coke to Bjorn" it first has to reason about the locations of the tagged object "coke" and person "Bjorn". It then forms and executes the plan to perform this task; namely to move to the kitchen, pick up the coke, move to the living room, where Bjorn is located, and hand the coke to Bjorn. Importantly, the logical operations of the robot is determined at a level of abstraction that reflects the natural intuitions of a human operator. In particular, locations in the house are defined in terms of named regions that reflect the everyday usage of a human operator, for example the different bedrooms, the kitchen, the living room, the study, the bathrooms, and the doorways separating these rooms.

Defining the different regions is currently undertaken by a human operator, however, as discussed in the section below on "Topological Reasoning", we are actively developing semi-automated means for the robot to be able to determine and label regions dynamically while interacting with the operator.



**Figure 1.** Sending the robot from the “master bedroom” to the “study” requires passing through the “dining room”, “living room”, and a number of doorways.



**Figure 2.** Segway RMP platform used for this project.

#### UNSW CSE Robotics Capabilities

The School of Computer Science and Engineering (CSE) at the University of New South Wales (UNSW) is one of the premier Australian computing schools. UNSW is the only Australian university ranked among the top 100 for Computer Science in the Academic Ranking of World Universities 2010 and obtained the most rankings for ICT disciplines in the recent Australian research Council (ARC) Excellence in Research Assessment (ERA) being ranked in five areas (including the area of this research) where other institutions were ranked in at most two areas. UNSW is a partner in the ARC Centre of Excellence for Autonomous Systems (CAS), which includes all investigators on this grant. With over 200 staff and research students, CAS is already the second largest robotics research group in the world with a leading reputation for both fundamental research and the application of this to

industry. In its first eight years, members of CAS produced over 800 research publications, gave 200 plenary/keynote/invited talks at international conferences, won 27 national and international prizes, and graduated 82 PhD students. All investigators on this grant are members of the UNSW node of CAS, which has an outstanding track record in international robotics competitions, including three times RoboCup world champion and three times runner-up in the four-legged league, 2nd place in the 2010 humanoid Aldebaran Nao-based Standard Platform League competition and 1st place in the technical challenges, and best autonomous robot award in 2009, 2010 and 2011 in the Rescue Robot league.

Autonomous Systems, an area in which this project directly falls, is a priority research area for UNSW. The University has invested more than \$1M in internal grants for building robotics research laboratories and purchasing equipment for research and research training. The UNSW node of CAS has several robot platforms including 5 ActivMedia Pioneers, 9 Aldebaran Nao humanoid robots, 3 urban search and rescue robots, 1 ActivMedia PeopleBot, 1 purpose built human-robot interaction robot, 1 Segway RMP, and several other purpose-built robot platforms that we are expecting to use in the course of our project. They are housed in three laboratories used by PhD students and undergraduate thesis students. In addition to these resources, the UNSW Faculty of Engineering has recently invested \$300K on a 1,152 core computing cluster that is available for large scale simulation and compute-intensive tasks that will be useful for this project.

The investigators have also established contact with some of the foremost research groups in this area throughout Australia and the world. In particular, the Computational Logic Laboratory at Simon Fraser University, the Computational Logic Group at the Technical University of Dresden, the Artificial Intelligence Research Group at the University of Freiburg, the Knowledge-based Systems Group at the Technical University of Aachen and the Knowledge Processing Laboratory at the University of Linköping.

This project has contributed to gaining the expertise within the UNSW CSE Robotics group for establishing the core capabilities necessary to tackle the Robocup@Home challenge. Investigator Thielscher has further contributed much of the funding for the purchase of equipment such as a SegwayRMP robotics platform with laser sensors, microphones and (soon to purchase) robot arms.

**Results and Discussion:** Describe significant experimental and/or theoretical research advances or findings and their significance to the field and what work may be performed in the future as a follow on project. Fellow researchers will be interested to know what impact this research has on your particular field of science.

This project has provided an exploration of the links between logical and non-logical systems as a means of achieving practical reasoning. Many of the main results have been theoretical to cement some of the basic science required to integrate both logical and non-logical forms of reasoning. The practical part of the project has been primarily undertaken within the context of challenges identified as part of the Robocup@Home competition. This competition seeks to focus the attention of the AI research community on easily identifiable problems for which the solution requires a multi-disciplinary approach.

The main outcomes of the project can be summarized as follows:

#### Reasoning with erroneous information

A formal framework establishing how a robot can sensibly respond to situations where it is given erroneous information by a human operator, but can only discover the error part-way through fulfilling a task. The motivation for this research has been to examine the realistic problems with which a robot will be confronted when operating within a home environment.

It further highlights the need for such a robot to develop common-sense notions in order to match the expectations of a human operator. The example scenario examined is based on a challenge from the Robocup@Home 2010 competition. A robot which is given the task of returning with the red cup from the kitchen table arrives in the kitchen to find no red cup but instead notices a blue cup and a red plate on the table, what should it do? The best course of action is to attempt to salvage the situation by relying on its preferences to return with one of the objects available.

Formally, this framework is developed in terms of the Situation Calculus extended with a notion of belief (Shapiro *et al.* 2011) by one of the investigators on this project. Preferences derived from the problem statement and common-sense knowledge are used to derive a plausibility ordering over all initial situations. It is this plausibility ordering that determines how the robot deals with the incorrect information; namely in this case that the robot chooses the blue cup over the red plate when it discovers that there is in fact no red cup.

While the Situation Calculus extended with beliefs provides an expressive formalism for tackling the problem of how an agent can recover when being given incorrect information, it does not provide for an efficient implementation of this formalism. Consequently, we adapt a recent extension of action logics with default reasoning (Baumann *et al.* 2010) developed by one of the investigators on this project and show how the problem can be transformed into a framework which can be efficiently implemented using Answer Set Programming (Gelfond 2008). The basic idea is to treat potentially erroneous information as something that is considered true by default but can be retracted should the agent make observations to the contrary.

The framework developed here provides an expressive and intuitive formalism for describing the problem of how a robot can dynamically respond to being given information that it later finds out to be incorrect. It further shows how this problem can be translated into a formalism that allows for efficient implementation, and therefore provides a path for implementation of this behaviour into existing robotics systems.

This work was presented and well received at the recent AAAI 2011 Workshop on Automated Action Planning for Autonomous Mobile Robots. Based on discussions at the workshop we are currently extending the work to show how our framework can encompass any PDDL description, as PDDL is a popular language used in the planning and robotics domains.

Maurice Pagnucco, David Rajaratnam, Michael Thielscher, and Hannes Strass, How to Plan When Being Deliberately Misled. *AAAI-11 Workshop on Automated Action Planning for Autonomous Mobile Robots (PAMR)*, 2011.

#### Approximate Reasoning

Examining aspects of propositional approximate logics. This work goes to the heart of determining logical satisfiability (SAT) within a resource-bounded reasoning context, and is therefore of practical interest in the implementation of reasoning for robotics. This work provides a logical characterisation of the well-known DPLL algorithm (Davis *et al.* 1962) when branching is restricted to a subset of the propositional variables (RDPLL). The work further establishes a link between this class of approximate logics and the notion of a *strong backdoor* of an unsatisfiable problem.

The DPLL algorithm forms the basis of a number of the efficiencies of modern SAT solvers. Despite recent advances in the practical efficiency of logical reasoners, determining propositional satisfiability remains the prototypical intractable problem (unless  $P = NP$ ). This computational intractability remains an important factor constraining the use of logical reasoners in strongly resource-bounded applications such as robotics. Fortunately, this need



to consider the limitations of real agents has lead to the examination of different sub-classes of propositional logic for which more computationally bounded behaviours can be determined. In particular the S-3 logics of Cadoli and Schaerf, best summarised in (Cadoli 1995), provide the first logical account of this form of approximate reasoning. These logics are parameterised by a set of propositional formula, the size of which determines the complexity of reasoning for the logic in question.

Unfortunately, the S-3 logics do not provide for a particularly accurate characterisation of the hardness of determining logical satisfiability. For example, a long chain of inferences requiring only the application of unit resolution is well-known to be solvable in polynomial time, yet would be characterised by the S-3 logics as being exponential in the size of the chain of inference. Consequently, in this paper we define an approximate logic for which this sort of inference can be shown to be solvable in polynomial time and therefore more closely characterises the inferences of real SAT solvers.

The family of logics developed in our research is defined in terms of a sub-class of Finger's logics of limited bivalence (Finger 2004) and shown to characterise the behaviour of the RDPLL algorithm. This work therefore opens the way for using more efficient solving techniques in resource-bounded applications such as robotics.

The work has been submitted and accepted at the 2011 Australasian Artificial Intelligence Conference and will be presented at the conference in December of this year.

David Rajaratnam and Maurice Pagnucco, From Approximate Clausal Reasoning to Problem Hardness. *Proceedings of the Twenty-Third Australasian Joint Conference of Artificial Intelligence*, to be published December 2011.

#### Topological Reasoning

Identifying the need to draw together the related but distinct areas of Qualitative Spatial Reasoning (QSR) and topological mapping in robotics. QSR is a well established AI sub-discipline looking at symbolic reasoning involving logical spaces. Possibly the most well-known axiomatisation of such a reasoning system is the Region Connection Calculus (RCC-8) (Randell *et al.* 1992). While providing for powerful spatial reasoning capabilities the precise means by which this can be used in robotics research is less certain. The main difficulty is that QSR systems are based on highly abstract notions of logical spaces and consequently do not translate easily to the coordinate systems that are the domain of robotics systems.


While there has been some work on providing a logical basis for topological map building within a robotics context, for example (Remolina 2004), much of the practical work in robotics is still focused on statistical methods for map construction. Map construction and path planning at this level consists of coordinate grids and being able to direct the robot to move to a specific grid location. Unfortunately, this topological view does not provide an appropriate level of abstraction to enable many of the types of reasoning tasks that would be important for a domestic robot. For example, a domestic robot would need to reason about locations such as rooms in a house, "the kitchen", "the bedroom", "the living room", as well as locations defined by sub-regions such as the "the table in the living room".

As part of this project, progress was made in bridging the gap between the data generated by robotic sensors and the needs of QSR reasoning systems. In particular we have developed a method for generating and updating a topological map based on the occupancy grid produced by a Simultaneous Localisation and Mapping (SLAM) algorithm. The method we develop draws on earlier work on region segregation (Thrun 1998) but applies a robust and more principled approach to delineating interesting regions.

An occupancy grid is a two dimensional representation of an area where each point on the

grid corresponds to the status of a region in space. A map point can be in one or two states: unknown (the area has not yet been explored), or a probability that the point in space is occupied by some obstacle. Typically a simple threshold is applied to this probability to make a decision about whether the region is free or occupied. The algorithm we have developed runs in real-time and builds the topological regions as the robot explores its environment and updates the occupancy grid.

SLAM and the construction of occupancy grids is one of the fundamental components of any robotics system. Consequently, our method for generating and updating topological maps can be applied on a wide variety of robotics platforms.




ion is based on identifying narrow spaces (red lines) that open into larger spaces. The Intersection between regions identify navigation points between these regions.

The basic motivation for this work is to develop robotics systems that can communicate and perform spatial reasoning at a cognitive level that has a natural correspondence to the abstract spatial notions of a human operator. For example, a human operator might instruct a domestic vacuum cleaner robot to clean “the kitchen”. The operator would expect the robot to clean an area that closely matches his/her notion of what constitutes the kitchen, without having to explicitly define a set of coordinates that provide the boundary points of the actual kitchen.

To enable the robot to perform spatial reasoning at this cognitive level we have developed a method whereby spatial regions are extracted from an occupancy grid and subsequently used to define human level abstract regions. At its most basic level the method is based on detecting narrow spaces that open out into larger spaces (Figure 3). These points are referred to as *critical corridor points*. These critical corridor points serve as the boundaries on which the regions are determined. A doorway into a room is the prototypical example of such a boundary, and can be used to define a room as a single region. However, doorways are not the only boundary points that can be used to define openings. For example, Figure 3 shows the result of the region detection algorithm when applied to the floor plan in Figure 1. While the main doorways are detected, a lounge chair is also used to separate the dining area from the living area, and the region detection algorithm is consequently able to distinguish these as distinct areas.





**Figure 4.** Regions are tagged and grouped to form a topological map of abstract

regions with a close correspondence to that of the human operator.

Once the basic regions have been extracted from the occupancy grid, the robot is able to name and combine these regions through interactions with a human operator. A typical use-case displaying this behaviour would be a new domestic robot following its owner around a home being told the names for different areas. Regions with the same name would be grouped (Figure 4) to more closely match the conceptual regions that the operator understands.

Providing the robot with a cognitive topological map allows for richer and more dynamic robot-human interactions. For example, by associating doorways with the intersection of regions a user could instruct a robot to “go to the kitchen via the hallway because the door between the dining room and kitchen is locked”. The robot could similarly display adaptive behaviour that can be understood by the operator; for example reporting back to a user that in fact “the door between the hallway and the kitchen is locked and not the door between the dining room and kitchen”.

While the current work has focus on reasoning about open regions, the same techniques can be applied to reasoning about occupied spaces. In this way relatively static objects in the home can be identified, such as the dining table. This would enable the robot to reason about a richer set of concepts. Furthermore, it can also be used by the robot in a more practical way to control its more expensive sensing actions. For example, when a robot is told to fetch a cup from the dining table it can reason that it only needs to apply computationally expensive image recognition techniques on the small occupied region identified as the dining table. Only if it were to fail to find the cup on the table would it need to consider looking more broadly around the dining room.

This work is currently being drafted into a paper for submission to the AAAI-12 special track on Robotics.

#### A New Programming Language for Practical Reasoning Systems

Logic programming is a powerful paradigm for programming autonomous agents in dynamic domains, as witnessed by existing languages such as Golog and Flux. We have developed a programming language that we called ALPprolog, an expressive, yet efficient, logic programming language for systems that have to reason about incomplete information and sensing actions. Our programming language is intended for the online control of agents, where actions are immediately executed. This starkly contrasts with offline reasoning, where agents may make assumptions to see where these are leading. ALPprolog was developed specifically for the efficient handling of large ground state representations, something that we consider to be practically useful. An implementation of ALPprolog is open source and can

be obtained at [alpprolog.sourceforge.net](http://alpprolog.sourceforge.net) (Drescher and Thielscher 2011).

## References

Ringo Baumann, Gerhard Brewka, Hannes Strass, **Michael Thielscher**, and Vadim Zaslavski, State Defaults and Ramifications in the Unifying Action Calculus. *Proceeding of the International Conference on Knowledge Representation and Reasoning*, 2010.

Marco Cadoli, Tractable reasoning in Artificial Intelligence, *LNAI 941*, Springer, 1995.

Martin Davis, George Logemann, and Donald Loveland, A Machine Program for Theorem-Proving, *Communications of the ACM*: 5(7), 394-397, 1962.

Conrad Drescher and **Michael Thielscher**, ALPprolog---A New Logic Programming Method for Dynamic Domains, Theory and Practice of Logic Programming (ICLP'11 Special Issue), **11**(4-5):451-468, 2011.

Marcelo Finger, Towards Polynomial Approximations of Full Propositional Logic. *In Proceedings of the Brazilian Symposium on Artificial Intelligence*, 11-20, 2004.

Michael Gelfond, Answer Sets, *In Handbook of Knowledge Representation*, 285-316, 2008.

Mohammed W. Kadous and **Claude Sammut**, MICA: Pervasive Middleware for Learning, Sharing and Talking, *In Proceedings of the IEEE PerCOM Workshops*, 2004.

Hector J. Levesque and **Maurice Pagnucco**, Legolog: Inexpensive Experiments in Cognitive Robotics, *In Proceedings of the Second International Cognitive Robotics Workshop*, Berlin, Germany, August 21-22, 2000.

Hector J. Levesque, Raymond Reiter, Yves Lesperance, Fangzhen Lin, Richard B. Scherl, GOLOG: A Logic Programming Language for Dynamic Domains, *Journal of Logic Programming*, 31:59-84, 1997.

David A. Randell, Zhan Cui, and Anthony G. Cohn, A Spatial Logic based on Regions and Connection. *In Proceedings of the International Conference on Knowledge Representation and Reasoning*, 1992.

Emelio Remolina and Benjamin Kuipers, Towards a General Theory of Topological Maps. *Artificial Intelligence* 152: 47-104, 2004.

Steven Shapiro, **Maurice Pagnucco**, Yves Lesperance, and Hector Levesque. Iterated Belief Change in the Situation Calculus. *Artificial Intelligence*, 144(1-2): 1-39, 2010.

Sebastian Thrun. Learning Maps for Indoor Mobile Robot Navigation. *Artificial Intelligence*, 99(1): 21-71, 1998.

**List of Publications and Significant Collaborations:** In standard format showing authors, title, journal, issue, pages, and date, please list the following:

- a) papers published in peer-reviewed journals,
- b) papers published in non-peer-reviewed journals or in conference proceedings,

Note: the following are peer-reviewed conferences and workshops.

Maurice Pagnucco, David Rajaratnam, Michael Thielscher, and Hannes Strass, How to Plan When Being Deliberately Misled. *AAAI-11 Workshop on Automated Action Planning for Autonomous Mobile Robots (PAMR)*, 2011.

David Rajaratnam and Maurice Pagnucco, From Approximate Clausal Reasoning to Problem Hardness. *Proceedings of the Twenty-Third Australasian Joint Conference of Artificial Intelligence*, to be published December 2011.

c) conference presentations,

Maurice Pagnucco, David Rajaratnam, Michael Thielscher, and Hannes Strass, How to Plan When Being Deliberately Misled. *AAAI-11 Workshop on Automated Action Planning for Autonomous Mobile Robots (PAMR)*, 2011.

d) manuscripts submitted but not yet published, and

e) manuscripts not yet submitted, and

Maurice Pagnucco, David Rajaratnam, Michael Thielscher, and Hannes Strass, Planning under Erroneous Information, proposed submission for KR-12 with extended version proposed for Artificial Intelligence Journal.

Maurice Pagnucco, David Rajaratnam, Michael Thielscher, and Claude Sammut, Practical Cognitive Spatial Reasoning for Domestic Robots, proposed submission for AAAI-12 special track on Robotics.

f) patents that resulted from this work.

Please also provide a list any interactions with industry or with Air Force Research Laboratory scientists or significant collaborations that resulted from this work.

**DD882:** As a separate document, please complete and sign the inventions disclosure form.

This document may be as long or as short as needed to give a fair account of the work performed during the period of performance. There will be variations depending on the scope of the work. As such, there is no length or formatting constraints for the final report. Keep in mind the amount of funding you received relative to the amount of effort you put into the report. For example, do not submit a \$300k report for \$50k worth of funding; likewise, do not submit a \$50k report for \$300k worth of funding. Include as many charts and figures as required to explain the work. A final report submission very similar to a full length journal article will be sufficient in most cases.